

International Business Machines Corporation Docket No.: YOR920040027US1
Harrington & Smith, LLP Docket No.: 909A.0156.U1(US)

Application for United States Letters Patent by:

James W. Cooper

Anna R. Coden

Stephen Boyer

**A SYSTEM AND METHOD FOR THE INDEXING OF
ORGANIC CHEMICAL STRUCTURES MINED FROM
TEXT DOCUMENTS**

**A SYSTEM AND METHOD FOR THE INDEXING OF ORGANIC CHEMICAL
STRUCTURES MINED FROM TEXT DOCUMENTS**

CROSS-REFERENCE TO A RELATED PATENT APPLICATION:

This patent application is related to copending and commonly assigned U.S. Patent Application S.N.

5 10/670,675, filed 09/24/2003, entitled: "System and Method for the Recognition of Organic
Chemical Names in Text Documents", by Anna R. Coden and James W. Cooper, the content of
which is incorporated by reference herein in its entirety.

TECHNICAL FIELD:

This invention relates in general to digital libraries and life science documents and, more
10 specifically, it relates to apparatus and methods for searching and analyzing scientific documents,
such as journal publications and patents, for the occurrence of names of organic chemicals and for
indexing their chemical structures.

BACKGROUND:

Regardless of the technology being used, most system for the analysis and indexing of documents
15 for search and information retrieval follow the same basic procedure. First the data are separated
into individual documents and each document is divided into text tokens. These tokens are then
combined into meaningful phrases and fragments that are indexed for retrieval. An index contains

data that is used for search and document analysis to process queries and identify relevant objects.

After the index is constructed, queries may be submitted to the search system. The query represents information that is desired by the user, and is expressed using a query language and syntax defined by the search system. The search system processes the query using the index data for the database

5 and a suitable similarity ranking algorithm. From this, the system returns a list of topically relevant objects, often referred to as a "hit-list". The user may then select relevant objects from the hit-list for viewing and processing.

In a network environment, the components of a text search system may be distributed across multiple computers. A network environment contains two or more computers connected by a local

10 or a wide area network, (e.g., Ethernet, Token Ring, the telephone network, and the Internet). A user accesses a hypermedia object database using a client application on the user's computer. The client application communicates with a search server (e.g., a hypermedia object database search system) on either the computer (e.g., the client) or another computer (e.g., one or more servers) on the network. To process queries, the search server needs to access just the database index, which may

15 be located on the same computer as the search server or on another computer on the network. The actual objects in the database may be located on any computer on the network.

A Web environment, such as the World Wide Web on the Internet, is a network environment where Web servers and browsers are used. Having gathered and indexed all of the documents available in the collection, the index can then be used, as described above, to search for documents in the

20 collection. Again, the index may be located independently of the objects, the client, and even the

search server. A hit-list, generated as the result of searching the index, will typically identify the locations and titles of the relevant documents in the collection, and the user then retrieves those documents directly using the user's Web browser..

Text mining of documents can also be performed as part of document indexing. Text mining
5 involves the recognition of document parts, such as paragraphs and sentences, and then the analysis
of each recognized document part (e.g., each sentence). Sentence analysis involves the tagging of
each word with its part of speech and then the parsing of each sentence into its component parts. The
result of sentence parsing is a parse tree of the parts and sub-parts of that sentence. This information
is typically stored in tables for retrieval. Frequently these tables are database tables with database
10 indexes associated with them.

Such parsing and data storage can then be used to deduce the overall meaning of the document and
the relations between parts of the document.

The ability to search patent and patent-related literature for information related to chemical entities
is particularly challenging. The nomenclature associated with chemical substances is difficult to
15 understand, and often inconsistent chemical terms are used to express the same or similar chemical
entities. Despite attempts to standardize the chemical nomenclature by international standards
committees such as the Union of Pure and Applied Chemist (IUPAC), these rules unfortunately have
not been consistently applied to chemical substances over time, particularly with respect to the
patent literature.

Historically, chemical entities were often referred to by "common names" and/or by inconsistently applied IUPAC rules. Often, terms that were acceptable in earlier years (for example 'potash') later gave way to other standards (potassium carbonate). Little or no effort has been made to "normalize" the chemical nomenclature of the intellectual property (IP) databases retroactively over the decades.

- 5 The problem of inconsistent naming is exemplified by considering the chemical names that have been applied to the drug ValiumTM (Valium is a registered trademark of Roche Products Inc.), the chemical structure of which is shown in Fig. 1. A list of some of the correct and incorrect names for ValiumTM that are found in the chemical and patent literature are shown in Table 1.

Table 1- Some of the chemical names used for ValiumTM in different databases

- 10 7-chloro-1-methyl-5-phenyl-2H-1,4-benzodiazepin-2-one
7-chloro-1-methyl-5-phenyl-3H-1,4-benzodiazepin-2(1H)-one
7-chloro-1-methyl-5-phenyl-1,3-dihydro-2H-1,4-benzodiazepin-2-one
7-chloro-1-methyl-2-oxo-5-phenyl-3H-1,4-benzodiazepine
1-methyl-5-phenyl-7-chloro-1,3-dihydro-2H-1,4-benzodiazepin-2-one
15 7-chloro-1,4-dihydro-1-methyl-5-phenyl-2H-1,4-benzodiazepin-2-one
7-chloro-1-methyl-5-3H-1,4-benzodiazepin-2(1H)-one

Additionally, in the case of pharmaceuticals, the names of compounds of interest often change over time as compounds become commercialized. This has led to the frequent use of trade names or generic names in the scientific literature or in medical databases, which are not reflected
20 retrospectively in the various IP databases. This has made it difficult to perform text searching for certain pharmaceuticals in the patent literature using commonly accepted phrases or definitions. For

example, one cannot simply type in the search term "aspirin" or "ValiumTM" into any of the IP databases and find the pertinent patents for those chemical substances. The problem is further exacerbated by the fact that different brand names are often used in different countries to address language considerations of the different geographical areas. In fact, there are as many as 149
5 different names that have been employed in the literature for the drug ValiumTM, a number of which are illustrated in Table 2.

Table 2 - Some of the trade names used to refer to ValiumTM

ALBORAL, ALISEUM, ALUPRAM, AMIPROL, ANSIOLIN, ANSIOLISINA, APAURIN,
APOZEPAM, ASSIVAL, ATENSINE, ATILEN, BIALZEPAM, CALMOCITENE, CALMPOSE,
10 CERCINE, CEREGULAR, CONDITION, DAP, DIACEPAN, DIAPAM, DIAZEMULS,
DIAZEPAM, DIAZETARD, DIENPAX, DIPAM, DIPEZONA, DOMALIUM, DUKSEN,
DUXEN, E-PAM, ERIDAN, EVACALM, FAUSTAN, FREUDAL, FRUSTAN, GIHTAN,
HORIZON, KIATRIUM, LA-III, LEMBROL, LEVUM, LIBERETAS, METHYL DIAZEPINONE,
MOROSAN, NEUROLYTRIL, NOAN, NSC-77518, PACITRAN, PARANTEN, PAXATE,
15 PAXEL, PLIDAN, QUETINIL, QUIATRIL, QUIEVITA, RELAMINAL, RELANIUM, RELAX,
RENBORIN, RO 5-2807, S.A.R.L., SAROMET, SEDAPAM, SEDIPAM, SEDUKSEN,
SEDUXEN, SERENACK, SERENAMIN, SERENZIN, SETONIL, SIBAZON, SONACON,
STESOLID, STESOLIN, TENSOPAM, TRANIMUL, TRANQDYN, TRANQUASE,
TRANQUIRIT, TRANQUO-TABLINEN, UMBRIUM, UNISEDIL, USEMPAX AP, VALEO,
20 VALITRAN, VALRELEASE, VATRAN, VELIUM, VIVAL, VIVOL, WY-3467

Additionally, many chemical and drug patents make use of Markush structure references. These

structures are generalized references to chemical structures where some substituent groups are specified in general terms, and a list of possible substitutents is enumerated. Thus, rather than a specific chemical compound being named, the Markush convention allows claimants to describe an entire series of compounds even if they have not specifically been synthesized or tested.

5

For example, and referring to Fig. 2, rather than representing toluene (methylbenzene) as $C_6H_5-CH_3$, the Markush formulation allows one to represent an entire series of substituted benzenes as C_6H_5-R , where R is, by convention, any of a large number of carbon chains of various sizes. This convention further increases the difficulty of locating a chemical compound by normal searching techniques.

- 10 In U.S. Patent No.: 6,304,869, Moore et al. describe a system to assign sub-structures to fragments given a complete structure connectivity description of a molecule, as well as a relational database system for storing this information. However, there is no concept of finding structures or substructures from names.

SUMMARY OF THE PREFERRED EMBODIMENTS

- 15 The foregoing and other problems are overcome, and other advantages are realized, in accordance with the presently preferred embodiments of these teachings.

This invention provides a system and a method to identify organic chemical nomenclature from text documents, and from that information to index chemical fragments and their structures and connectivity. This process can involve the grouping of multi-word entities into a single logical

entity, and then the parsing of that entity into names of substructures. The text documents can be either well edited (where the rules for denoting such entities are followed) or ill formed. The system and method in accordance with this invention may be applied to both types of documents. Furthermore, only relatively small dictionaries need to be used.

- 5 Disclosed is a method, a computer program product and a system for processing documents that contain chemical names. In a system embodiment the system can include one computer, or a plurality of computers at least two of which are coupled together through a data communications network. The system has a unit to parse document text to recognize chemical name fragments; a unit to recognize any substructures present in the chemical name fragments; and a unit to determine
- 10 structural connectivity information of the chemical name fragments and recognized substructures and to store the determined structural connectivity information in a searchable index.

The determined structural connectivity information is preferably stored in a searchable structure index, and the system further includes a unit to store text associated with processed documents in a text index, and a unit to search the text index using at least one of a fragment name and a

- 15 substructure name and to search the structure index by at least one of fragment connectivity and substructure connectivity. At an intersection of the search results from the structure index and the text index, the system operates to identify at least one document that contains a reference to a corresponding chemical compound.

The unit that determines structural connectivity information looks up recognized fragments and

- 20 substructures in a structure dictionary. In the preferred embodiment the structure dictionary is at

least one of a MOL dictionary and a SMILES dictionary.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other aspects of these teachings are made more evident in the following Detailed Description of the Preferred Embodiments, when read in conjunction with the attached Drawing

5 Figures, wherein:

Fig. 1 shows the chemical structure of ValiumTM;

Fig. 2 illustrates a conventional Markush formulation, where Fig. 2A illustrates the chemical composition of toluene, and where Fig. 2B illustrates a Markush representation that includes toluene;

10 Fig. 3 shows various chemical substructures parsed from the chemical name for ValiumTM;

Fig. 4 shows a MOL file representation of a diazepin fragment;

Fig. 5 illustrates correlating chemical structure fragments with connectivity tables;

Fig. 6A is a logic flow diagram that illustrates the overall flow of an indexing algorithm in accordance with this invention;

Fig. 6B illustrates the logic flow of a presently preferred search algorithm; and

Fig. 7 is a block diagram of an exemplary embodiment of a computer system that is suitable for practicing the method of this invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

5 By way of introduction, this invention uses a series of regular expressions, rules, and two small dictionaries to recognize chemical name fragments and combine them into organic chemical names. The use of the system and method of this invention is valuable in assisting parsers in recognizing multi-word chemical names that might otherwise be recognized as small fragments separated by punctuation that is part of the chemical names. Then, each chemical name is decomposed into
10 fragments and indexed for text searching. If the fragment is known from a dictionary of known chemical structure fragments, the connectivity of this substructure is saved for indexing as well.

In more detail, it may be first assumed that the algorithm described in the above-referenced commonly assigned U.S. Patent Application S.N. 10/670,675 has identified as a chemical compound the string:

15 7-chloro-1-methyl-5-phenyl-3-dihydro-2H-1,4-benzodiazepin-2-one

The system and method in accordance with this invention then parses the above-given string into component fragment names, and indexes each of them separately. In this elementary example, the

system and method finds the fragments between the hyphens: i.e., chloro; methyl; phenyl; 2H; dihydro; benzodiazepin; one; and produces a candidate list.

The candidate list is filtered in several steps, using several pattern rules and a small dictionary of known chemical substructures, resulting in the following list: chloro; methyl; phenyl; benzo; diazepin; one. Note that in this context the "one" fragment is not a number, but refers to a ketone substructure. The structures corresponding to these extracted fragment names are shown in Fig. 3.

As a result, the system and method are thus enabled to associate these name fragments with a dictionary of chemical structure fragments (CSFs) in an efficient manner. In addition, the dictionary of CSFs can contain graphical descriptions and be used provide a visual display to enhance the overall search process for compounds containing those entities.

Conversion of the chemical name fragments (CNFs) to CSFs implies that irrespective of the name a particular researcher or searcher uses, the user can search for any of these fragments by structure without having to select or specify the actual name used in the document.

Thus, while the numerous variations in the name of Valium™ in Tables 1 and 2 are too extensive for a text search to be helpful, a search for the fragments by structure is much more likely to be successful.

In mining information from text documents, such as patents and technical articles, it is critical that long multi-word organic chemical nomenclatures be recognized properly so they can be grouped as

single logical entities and correctly indexed. In the above-referenced commonly assigned U.S. Patent Application S.N. 10/670,675 the inventors Coden and Cooper previously described a system and method for grouping such nomenclature into logical entities without the need to provide large chemical dictionaries. This invention makes use of a search engine, such as one known as a 5 JuruXML™ search engine available from the assignee of this patent application, and a table of substructure names and connectivity. Such a table could, for example, be stored in a relational database such as one known as DB2™, also available from the assignee of this patent application.

Organic chemical names can be long, complex and may consist of several words separated by spaces. Organic chemical names should be recognized as a single noun phrase in order for the 10 parsing of sentences in technical documents to proceed effectively. For example, terms such as chloroacetic acid, 4-allyl-2,6-dimethylphenol, 5-aminoalkyl-pyrazolo[4,3-D]-pyrimidine and 4-nitrobenzyl chloroformate each present specific term recognition challenges that previously could only be resolved by reference to a multi-million word chemical dictionary.

Further, while there are specific chemical rules for the spelling, spacing and punctuation of such 15 chemical entities, they are not always rigorously followed, especially in the patent literature. Examples abound of chemical names broken up by incorrect spaces or hyphens which must be recombined for the overall term to be recognized successfully.

There are several common methods of representing the connectivity of organic chemical structures. Two such formats are referred to as MOL files and SMILES files. MOL files (from Molecular 20 Design Ltd) contain the coordinates of each atom along with a connectivity matrix, while SMILES

files represent connectivity using letters for each atom and symbols for the various bonds, see "SMILES 1. Introduction and Encoding Rules", Weininger, D., *J.Chem. Inf. Comput. Sci.* 1988, 28,31. Both the MOL and SMILES approaches can be used to represent extremely complex structures. This invention assumes that a table of such connectivity representations is available for
5 the common molecular fragments that can be named in chemical structures, such as methyl, phenyl and so forth, in at least one of these common formats, or that a parser exists to read data in these formats and convert it into an internal structural representation. Such parsers are well known in the art and are readily available.

This invention generally has two phases or aspects: a first relates to the indexing of chemical
10 structural fragments, and a second relates to returning query results of such fragments. In the indexing phase of the invention it is assumed that a series of chemically-related documents, such as chemical or drug patents or articles, are scanned and indexed.

Indexing Patents and Articles

In the indexing phase, each document is analyzed and the text indexed by a search engine. Then,
15 organic chemical names are identified and the fragment names in these names are also added to the index. Finally, for each fragment that the analysis system finds within each organic chemical name, it looks up that name in a substructure dictionary. If that substructure is found in the dictionary, it is added to a structure index for that document.

Each chemical name is broken into fragments using a tokenizer that separates tokens based on any
20 of the punctuation characters -()#0123456789 and space. Then, those fragments are eliminated

that contain numbers.

For example, for the chemical name:

7-chloro-1-methyl-5-phenyl-3-dihydro-2H-1,4-benzodiazepin-2-one

that was mentioned above, the parser extracts the substructure fragments: chloro; methyl; phenyl;

5 dihydro; benzodiazepin; one. Next a substructure string search is applied to these fragments, breaking them down further: chloro; methyl; phenyl; dihydro; benzo; diazepin; one. These strings are then looked up in a structure dictionary, and for those that are found, a substructure entry is made. The corresponding SMILES strings are as follows:

Cl

10 [CH3]

c1ccccc1

c1ccccc1

c1ncccncc1

C=O.

15 An analogous set of entries can also be made for the MOL file representations, which represent atomic coordinates and connectivity numerically. For example, the MOL file representation of the diazepin fragment is shown in Fig. 4.

It should be noted that this approach goes well beyond just synonym expansion, as the method expands molecule names to their substructures and represents these substructures so that they can

20 be searched for without reference to the name used in that particular molecular name.

Searching the Indexes

In the search phase, the user enters search terms and structures. For example, to enter a SMILES format substructure query, the user would enter "c1ccccc1" for a phenyl group and "C=O" for a carbonyl group, along with one or more search terms such as "antidepressant" or "antibiotic" (or 5 whatever other term(s) that may have been saved in the text index 715 shown in Fig. 7, such as author name, article title or, for the case of a patent document, the patent number, the assignee name, etc.) The results of the structure search and the text search are combined and their intersection is returned to the user.

In commonly assigned U.S. Patent Application S.N. 10/670,675 there is described a system and a 10 method for recognizing chemical names algorithmically, without resort to large compendia of chemical knowledge. Described herein is a system and method for indexing chemical names into chemical fragments which can be correlated with chemical connectivity tables.

The method includes recognizing the chemical name, and finding its fragmentary components. The 15 fragments are indexed for insertion into chemical connectivity tables (such as in MOL and/or SMILES representation) and possibly also for text search, and those fragments whose substructures are known are indexed into chemical connectivity tables as well. Furthermore, the method handles those chemical fragments that were written not following the standard rules of writing such entities, or that contain erroneous spaces and/or characters caused by, for instance, the use of OCR software.

Recognizing Organic Chemical Fragments

20 Algorithms used for indexing organic chemical names are now described. The use of the system and

method results in both a text-based search index by chemical fragment name and a substructure search based on chemical structure connectivity.

- A. It is assumed that each document to be indexed is parsed using text analytics and organic chemical name recognition. A presently preferred technique for performing this function is 5 described in the commonly assigned U.S. Patent Application S.N. 10/670,675, however the use of this invention is not limited only to the use of only the technique described in commonly assigned U.S. Patent Application S.N. 10/670,675.
- B. Each organic chemical name is broken into chemical subtokens wherever parts of the longer name are separated by specific tokens. In this preferred embodiment, these tokens include, but need 10 not be limited to, hyphens, parentheses, brackets and braces.
- C. Each subtoken that does not contain a number is added to a search index.
- D. Each subtoken that does not contain a number is looked up in a chemical fragment dictionary and its connectivity information retrieved, if it exists. This connectivity information is added to a chemical substructure index (shown in Fig. 7 as the structure index 719).
- 15 E. The chemical substructure index 719 may be, as non-limiting examples, a text file, an XML file, or a relational database. Each substructure may represent connectivity in either the MOL file or SMILES representation.

F. The user of a search system enters one or more fragment names and selects one or more substructures, either by name or by graphical representation (such as by using a pointing device to select a particular graphical representation of a chemical structure from a pull-down or pop-up menu of possible choices).

5 G. The search system returns the identification of documents, and possibly copies of the documents themselves, where chemical compounds have been found that contain the selected substructure names and the connectivity specified by the selected substructures and/or fragments.

Fig. 5 shows the general process of correlating chemical structure fragments with connectivity tables, where CTF represents chemical connection table fragments. The CNF represents chemical
10 name fragment and the CSF is the chemical structure fragment, as defined above.

Fig. 6A describes indexing a collection of documents. Each document is read in from a file (block 600) and indexed (block 601) in a conventional manner using a search engine, such as the JuruXML™ search engine. In the presently preferred embodiment the algorithm described in the commonly assigned U.S. Patent Application S.N. 10/670,675 is then used to identify organic
15 chemical names (block 602). Each organic chemical name is separated into sub-tokens, separated by, for example, hyphens, spaces and parenthesis (block 603).

In a loop, the system and method tests to see if more fragment tokens remain (block 604). If they do remain, the next fragment token is obtained and tested to see if it occurs in a dictionary of SMILES fragments (block 605). If it does, the SMILES expression is added to the structure index

(block 606). Then, a test is preferably also made to determine if the fragment occurs in the MOL file dictionary (block 607). If it does, it is added to the structure index as well (block 608). A test is then made to determine if there are more documents to be processed (block 609). If there are, control passes to block 600 to continue processing the remaining documents, and if not the indexing 5 operation is completed. The processing of fragments implies as well the processing of substructures that may make up a certain fragment.

Fig. 6B describes how a user performs a search against the system. The user first enter search terms (block 610), and then optionally the user enters substructures either as text, such as SMILES text (block 611); or using a graphical user interface (GUI) as a pointer to a graphical list of structures 10 (block 612), or the user enters a pointer to a MOL file (block 613). The system then searches the text index (block 614) and the structure index (block 615), and what is returned is the intersection of the results (block 616). The results identify one or more documents where there are found chemical compounds that contain the desired substructures and the specified search terms.

Fig. 7 is a block diagram of a document processing system 700 that incorporates functional blocks 15 (702-714) of the commonly assigned U.S. Patent Application S.N. 10/670,675, and that further incorporates functional blocks (715-721) in accordance with this invention. As such, it is assumed that the system 700 contains a standard tokenizer 702 for separating input document text 704 into tokens 706 based on blank spaces. The tokens 706 are examined where they match a set of defined patterns. Furthermore the tokens 706 are examined in the context of the adjacent tokens, to 20 determine whether the tokens 706 are part of a chemical fragment. More specifically, the system 700 includes a token processing unit 705 for determining semantic meaning of words, such as by

assigning corresponding associated parts of speech to words found in the document. The token processing unit 705 may be constructed to include sub-units 707, 709 and 711 for applying a plurality of regular expressions and rules, and a plurality of dictionaries, to recognize organic chemical name fragments (sub-unit 707), for combining recognized organic chemical name
5 fragments into a complete organic chemical name (sub-unit 709), and for assigning the complete organic chemical name with one part of speech, preferably a noun (sub-unit 711). The aforementioned dictionaries of the system 700 can include a prefix dictionary 708 (containing a list of common prefixes for the technical terms of interest), a suffix dictionary 710 (containing a list of common suffixes for the technical terms of interest), and an optional negative dictionary 712. The
10 negative dictionary 712, if used, contains words that may occur within the input document text 704, but that do not form a meaningful part of a technical term (e.g., do not form a part of an organic chemical compound). Basically, the negative dictionary 712 includes a list of words that can be ignored. Examples of words that may be found in the negative dictionary 712 are "saline" and "formula". It should be appreciated that the contents of the dictionaries 708, 710 and 712 can change
15 and evolve over time, and over the use of the system 700, either manually or automatically. The plurality of regular expressions (patterns) and rules can be stored in a database 713, and may also change and evolve over time, and over the use of the system 700, either manually or automatically. The output of the token processing unit 705 can form an input to a further unit 714 that parses sentences into their component parts based at least in part on the determined semantic content, such
20 as assigned parts of speech (including the noun part of speech assigned to recognized organic chemical names).

In the operation of the token processing unit 705 the application of regular expressions and rules

results in punctuation characters being one of maintained or removed between chemical name fragments as a function of context. The regular expressions can include a plurality of patterns, where individual patterns can be at least one of characters, numbers and punctuation. For example, the punctuation can include at least one of parenthesis, square bracket, hyphen, colon and semi-colon,
5 and the characters can include at least one of upper case C, O, R, N and H, as well as strings of at least one of lower case xy, ene, ine, yl, ane and oic.

The output of the sentence parser 714 is applied to a text index 715, where for each examined document there is a list of associated text found in that document. Examples of text can include author(s) names and keywords such as "antidepressant", "antiseptic", "protein", etc., as well as the
10 recognized names of chemical compounds.

In addition, the output of the fragment recognition block 707 is applied to a substructure recognition block 716. Using again the example given above of the organic chemical name:
7-chloro-1-methyl-5-phenyl-3-dihydro-2H-1,4-benzodiazepin-2-one,
the fragment recognition block 707 extracts the substructure fragments: chloro; methyl; phenyl;
15 dihydro; benzodiazepin; one. These fragments are then applied to the substructure recognition block 716, where the substructure string search is applied to these fragments, breaking them down further where possible into the substructures: chloro; methyl; phenyl; dihydro; benzo; diazepin; one. These substructure strings (some being the original fragments, and some possibly being substructures that make up one or more of the fragments) are then input to a substructure lookup block 717 where, in
20 cooperation with at least one structure dictionary 718 (e.g., one or both of a MOL or SMILES dictionary), the substructure strings are looked up in the structure dictionary 718, and for those that

are found, a substructure entry is made in a structure index 719.

Coupled with the structure index 719 and with the text index 715 is a searcher 720 and a user interface (UI) 721, such as a graphical user interface (GUI) comprised of a display 721A and a keyboard 721B. By means of the UI 721 and searcher 720 the user is enabled to perform the search
5 method disclosed above in association with Fig. 6B.

The foregoing description has provided by way of exemplary and non-limiting examples a full and informative description of the best method and apparatus presently contemplated by the inventors for carrying out the invention. However, various modifications and adaptations may become apparent to those skilled in the relevant arts in view of the foregoing description, when read in
10 conjunction with the accompanying drawings and the appended claims. For example, only one of the MOL or SMILES chemical representation systems may be used, or another type of representation system may be employed alone or in combination with one or both of the MOL and SMILES systems. However, all such and similar modifications of the teachings of this invention will still fall within the scope of this invention.

15 It should be further appreciated that the system 700 could be implemented in a network environment, and that components of the system 700 may be distributed across multiple computers. The network environment may contain two or more computers connected by a local or a wide area network, (e.g., Ethernet, Token Ring, the telephone network, and the Internet), and a user may access a hypermedia or other object database using a client application on the user's computer. The
20 client application may communicate with a search server (e.g., a hypermedia object database search

system) located on a client computer or another computer (e.g., one or more servers) on the network. To process queries from users, the search server may access a database index, which may be located on the same computer as the search server or on another computer on the network. The document objects in a database may be located on any computer on the network. In this invention certain of 5 the functional units and modules shown in Fig. 7, such as the token processing unit 705 and components of the token processing unit 705, as well as the substructure recognition 716, substructure lookup, structure dictionary 718 and structure index 719 (as well as the text index 715) may be located on two or more computers, and may be coupled together by one or more data communications networks. One or more of the connections between the tokenizer 702 and the token 10 processing unit 705, and/or the token processing unit 705 and the sentence parser 714, and/or the fragment recognition unit 707 and the substructure recognition unit 716, may also be implemented over data communications networks, including local and wide area networks, such as the Internet. The input to the tokenizer 702 and the output from the sentence parser 714 may also be implemented using one or more networks. The user may query the system 700 over a network, such as the 15 Internet, and the system 700 may form a part of a network-based, e.g., a Web-based, service such as, by example only, a data mining type of service.

Further, while the method and apparatus described herein are provided with a certain degree of specificity, the present invention could be implemented with either greater or lesser specificity, depending on the needs of the user.

20 Further still, some of the features of the present invention could be used to advantage without the corresponding use of other features. As such, the foregoing description should be considered as

merely illustrative of the principles of the present invention, and not in limitation thereof.